4 Example Demonstration Grids

STAR/HRM achieves robust, effective, Terabyte-scale multi-file replication

Members of the STAR experiment and the Scientific Data management group at Lawrence Berkeley National Laboratory collaborated to automate data transport between the RHIC Computing Facility (RCF) storage system at Brookhaven National Laboratory, and the NERSC storage system at LBNL, using Hierarchical Resource Managers (HRM) software developed by the Physics Grid Projects.

Since STAR began data taking two years ago, 10's of terabytes (TB) have been transferred at rates of about 1 TB/week with considerable effort and ad hoc methods. Using the grid has boosted implementation rates up to 8 MB/sec for the WAN stage, using the GridFTP data transport program.

After resolving some end-point configuration issues we expect that rates of 3-4 TB/week will be easily achieved during STAR's 2003 data-taking run. The group has defined a standard interface to storage systems, and interoperable versions for ENSTORE at Fermilab and JASMINE at Thomas Jefferson National Laboratory have been developed as part of the collaboration.

HRMs: High Level View

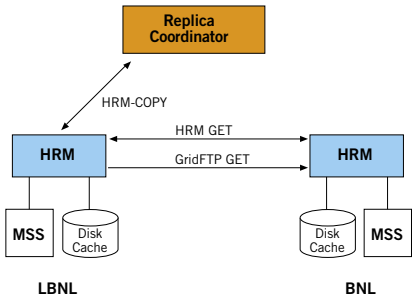

US/CMS Simulates Physics on the Data Grid

US/CMS has generated its first grid-enabled, production-quality, simulated data on a grid. The CMS Monte Carlo Production (CMS-MOP) scripts use VDT grid services from Condor-G, DAGMAN, GDMP, and the Globus Toolkit package. The CMS-MOP distributed production system employs a tier-like hierarchy in which a production manager at a Tier-1 center distributes production jobs to several remote Tier-2 sites. Once generated at the Tier-2 sites, the simulated data is automatically published back to the Tier-1 center as well as replicated to selected Tier-2 sites.

More than 150,000 proton-proton collision events inside the CMS detector have been simulated using CMS-MOP. Plans are ramping up for 1.5M events to be delivered to the experiment's physicists by the end of 2002, to help in designing the detector.

Data Grid Hierarchy (CMS)

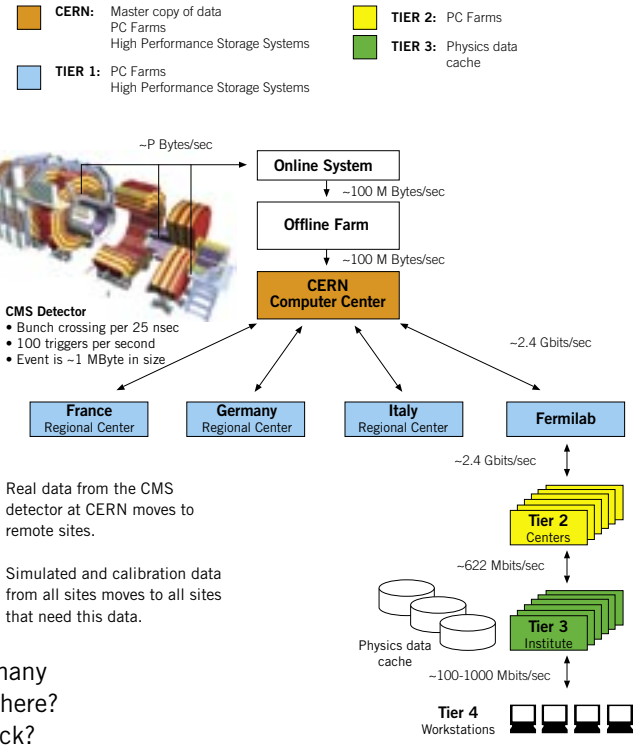

Real data from the CMS detector at CERN moves to remote sites.

Simulated and calibration data from all sites moves to all sites that need this data.

Real World Issues

The USCMS integration exercise has demonstrated that Grids still present significant challenges in harnessing distributed resources. Issues of data and security had to be overcome: How do you get your software and data to many remote systems and know it's there? How do you get your results back? How do you allow people on the Grid to trust you? Issues of heterogeneity and error recovery had to be addressed: To use other sites' resources, you need to interface with many batch systems; the Grid means more errors, more crashes, more mysterious failures. Unanticipated errors came up: key machines crashing in the middle of a run; Grid credentials expiring in the middle of a run; jobs completing successfully but their results being lost before being sent back; various pieces of middleware performing unpredictably; the network going down.

Applications of the Virtual Data Toolkit

The virtual data toolkit has been used with great success to study many promising science applications of the virtual data paradigm and grid computing. The Trillium collaborations have been working to adapt virtual data mechanisms and grid infrastructures built from the VDT, and apply them to scientific challenges, such as high-energy physics simulations, and the astronomy problem of identifying clusters of galaxies.

CMS Simulations

In high-energy physics, the VDT was applied to high-volume Monte Carlo simulation of the complex, data-intensive particle collisions that will take place in the CMS detector of the Large Hadron Collider. CMS physicists have demonstrated the use of virtual data definitions and requests to create large numbers of simulated events from a simple virtual data system request. At the University of Florida, the application steps of the complex multi-stage pipeline were defined in a virtual data catalog; then the derivations required for the simulation were entered into the catalog with a simple script, and the workflow needed to produce the data was generated by the virtual data system. This technique has been tested at the scale of 150,000 events, and much larger efforts to generate real usable physics simulation data are under way.

SDSS Cluster Finding

Trillium collaborators in the Sloan Digital Sky Survey and GriPhyN project have been using the classical astronomy problem of galaxy cluster-finding as a test case to study the value of virtual data techniques for managing and analyzing astronomy survey data. The techniques developed so far have provided SDSS scientists with the ability to trace the derivation of their data through virtual data mechanisms, and to re-execute portions of those derivations as algorithms and/or source data evolve. Visual interfaces linked to online SDSS data repositories permit navigation through online SDSS survey images, and tie interactive visual exploration of clusters with detailed information about the processing chains leading to each cluster identification.

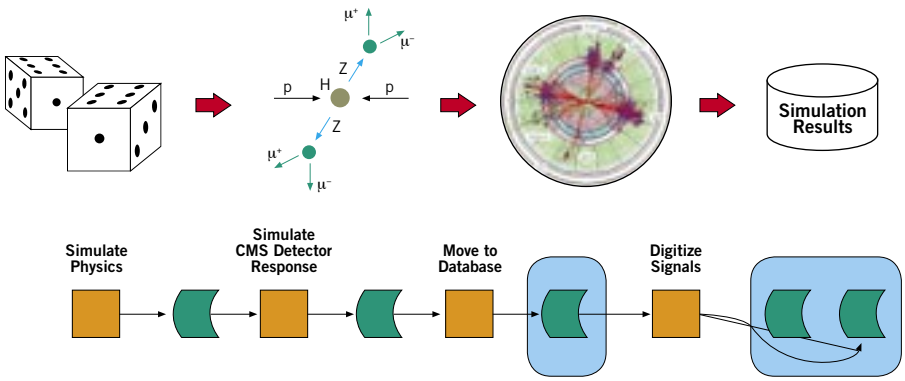

BaBar Data Streaming

Babar's data rates have outgrown the local computing center's capability to provide sufficient analysis for the physicists. Data is steadily streamed across the Atlantic to peer sites in France, England and Italy with current daily rate of over 150 Mbits/scs each on Internet2 and Esnet. If this rate of growth continues, in 10 years' time we will need 24 Terabits/sec for BaBar alone.